

SUPPLEMENT TO “PARAMETRIC INFERENCE AND DYNAMIC
STATE RECOVERY FROM OPTION PANELS”
(*Econometrica*, Vol. 83, No. 3, May 2015, 1081–1145)

BY TORBEN G. ANDERSEN, NICOLA FUSARI, AND VIKTOR TODOROV

In this supplement, we present an additional Monte Carlo exercise and additional diagnostics for the empirical application in the paper as well as estimation results for alternative stochastic volatility model specifications. We also provide further details regarding the computations in the empirical part and the Monte Carlo study.

S1. ADDITIONAL MONTE CARLO EVIDENCE

IN THIS SECTION, we report our findings for the performance of our developed estimator and tests on simulated data from the following two-factor stochastic volatility model for the underlying stock price X under the *risk-neutral distribution*:

$$(72) \quad \begin{aligned} \frac{dX_t}{X_{t-}} &= \sqrt{V_{1,t}} dW_{1,t} + \sqrt{V_{2,t}} dW_{2,t}, \\ dV_{1,t} &= \kappa_1(\bar{v}_1 - V_{1,t}) dt + \sigma_1 \sqrt{V_{1,t}} dB_{1,t}, \\ dV_{2,t} &= \kappa_2(\bar{v}_2 - V_{2,t}) dt + \sigma_2 \sqrt{V_{2,t}} dB_{2,t}, \end{aligned}$$

where $(W_{1,t}, W_{2,t}, B_{1,t}, B_{2,t})$ is four-dimensional Brownian motion with correlations $\rho_1 = \text{corr}(B_{1,t}, W_{1,t})$ and $\rho_2 = \text{corr}(B_{2,t}, W_{2,t})$. The parameter vector is given by $\theta = (\rho_1, \bar{v}_1, \kappa_1, \sigma_1, \rho_2, \bar{v}_2, \kappa_2, \sigma_2)$ and the parameter values used in the Monte Carlo are reported in Table VII. The observation setting is exactly the same as in the Monte Carlo study reported in Section 5 of the paper with the only exception being that now the moneyness range is reduced to $[-2, 1] \cdot \sigma\sqrt{\tau}$ because, for the model (72) without any jump component, the OTM option prices in the range $[-4, -2] \cdot \sigma\sqrt{\tau}$ are very close to zero, and not representative of what is observed in practice.

The precision in recovering the parameters is reported in Table VIII. Overall, the parameters are estimated quite well and the biases are close to negligible.

Turning next to the diagnostic tests, Table IX reports on the size of the various tests developed in Section 4 of the paper. Generally, the small sample behavior is satisfactory. The tests for the fit to the option panel are almost perfectly sized, with only mild over-rejection for the OTM short-maturity puts. The omnibus test for the parameter stability test under-rejects slightly while the volatility test rejects a bit too frequently.

Finally, in Table X we report on the tests for stability of individual parameters. They perform well, except for \bar{v}_1 and \bar{v}_2 , where we notice a somewhat

TABLE VII
PARAMETER SETTING FOR THE NUMERICAL EXPERIMENTS

Under \mathbb{P}				Under \mathbb{Q}			
Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
ρ_1	-0.500	ρ_2	0.000	ρ_1	-0.500	ρ_2	0.000
\bar{v}_1	0.010	\bar{v}_2	0.015	\bar{v}_1	0.020	\bar{v}_2	0.020
κ_1	4.000	κ_2	30.000	κ_1	2.000	κ_2	15.000
σ_1	0.250	σ_2	0.700	σ_1	0.250	σ_2	0.700

TABLE VIII
MONTE CARLO RESULTS: ESTIMATION OF THE RISK-NEUTRAL PARAMETERS

Parameter	True Value	Median	IQR	Parameter	True Value	Median	IQR
ρ_1	-0.500	-0.502	0.116	ρ_2	0.000	0.000	0.036
\bar{v}_1	0.020	0.020	0.001	\bar{v}_2	0.020	0.020	0.002
κ_1	2.00	1.993	0.242	κ_2	15.000	14.938	1.971
σ_1	0.250	0.249	0.043	σ_2	0.700	0.699	0.089

TABLE IX
MONTE CARLO RESULTS: DIAGNOSTIC TESTS^a

Test	Nominal Size of Test		
	1%	5%	10%
Panel A: Fit to Option Panel			
Out-of-the-money, short-maturity puts	3.1%	7.8%	12.7%
Out-of-the-money, short-maturity calls	1.0%	4.0%	7.6%
Out-of-the-money, long-maturity puts	1.0%	4.9%	10.0%
Out-of-the-money, long-maturity calls	0.5%	4.0%	9.1%
Panel B: Parameter Stability			
	1.2%	2.6%	5.4%
Panel C: Distance Implied-Nonparametric Volatility			
	3.6%	9.4%	14.9%

^aTable description as for Table IV in the paper.

TABLE X
MONTE CARLO RESULTS: TESTS FOR STABILITY OF INDIVIDUAL PARAMETERS^a

Parameter	Nominal Size			Parameter	Nominal Size		
	1%	5%	10%		1%	5%	10%
ρ_1	2.4%	6.0%	12.6%	ρ_2	1.2%	3.2%	7.4%
\bar{v}_1	0.2%	1.2%	2.0%	\bar{v}_2	0.4%	1.2%	2.2%
κ_1	1.4%	5.2%	10.2%	κ_2	2.2%	4.8%	10.0%
σ_1	1.6%	5.2%	12.2%	σ_2	1.0%	3.4%	7.8%

^aThe parameter stability test is given in equation (18) in the paper.

larger degree of under-rejection (particularly at the 10% level). As seen from Table VIII, these two parameters are recovered extremely precisely, so the under-rejection stems from a slight over-estimation of their asymptotic variances.

Overall, the simulation evidence confirms that our inference technique works satisfactorily even in the more challenging case when the asset dynamics is governed by a two-factor stochastic volatility model.

S2. FURTHER DETAILS ON THE EMPIRICAL APPLICATION

S2.1. *Additional Diagnostics for the Empirical Application*

We now present additional diagnostic tests for the two estimated models (1) and (22). First, in Table XI, we report the rejection rates for the formal tests for fit to the different regions of the option surface as well as the tests for equality of the nonparametric and the option pricing model implied spot volatility estimates.

Panel A of Table XI shows also that moving from the one-factor to the three-factor model (22) provides a near uniform improvement in the model's ability to fit the different parts of the option surface over the sample period. The improvement is very significant for the short- and long-maturity OTM puts and ATM options. There is only a small increase in the rejection rates for the long-maturity OTM calls. Overall, the three-factor model improves the average fit of the one-factor model, measured in terms of the root-mean-squared-error, by almost 50%.

Next, Table XII reports rejection rates for pairwise tests of individual parameter stability across each calendar year within our sample. The table reveals significant variation in the parameter estimates, but also a dramatic improvement in the stability for some parameters as we move from the one-factor to the three-factor model, suggesting improved model specification. Nevertheless, some parameters in the three-factor model are quite unstable, most no-

TABLE XI
DIAGNOSTIC TESTS FOR S&P 500 OPTION DATA^a

Test	One-Factor Model Nominal Size		Three-Factor Model Nominal Size	
	1%	5%	1%	5%
Panel A: Fit to the Option Panel				
OTM, short-maturity puts	45.92%	60.92%	18.03%	35.00%
ATM, short-maturity puts	70.00%	77.89%	16.05%	30.39%
OTM, short-maturity calls	30.39%	49.08%	20.66%	37.11%
OTM, long-maturity puts	60.39%	72.89%	16.71%	30.53%
ATM, long-maturity puts	62.76%	74.34%	18.42%	32.37%
OTM, long-maturity calls	9.08%	15.53%	14.34%	34.87%
Panel B: Root-Mean-Squared-Error of IV Option Fit				
	3.06%		1.59%	
Panel C: Equality of Implied and Nonparametric Volatility				
	53.29%	63.42%	51.05%	62.63%

^aPanel A reports rejection frequencies across the full sample for the option fit to specific regions of the option surface at the end of trading on Wednesdays. This test relies on Corollary 1, using the first two maturities for the first three tests and all remaining options with maturity less than one year for the last three. OTM puts and calls, ATM options, and short- versus long-maturity options are defined in Figure 3. Panel B provides the root-mean-squared-error of the model-implied BSIV relative to the market mid-quote BSIV across all options used during estimation over the full sample. The test in Panel C is defined in Corollary 3.

tably the persistence parameters κ_1 and κ_2 . Again, it is evident that neither model is correctly specified. In fact, for the two models, the joint test for stability of the full parameter vector across any two consecutive years has a 100% rejection rate. If anything, this confirms the power of our tests and reiterates the point that none of the models provide an ideal fit to the complex option surface dynamics.

Finally, Figure 7 depicts the nonparametric and the option-implied volatility series extracted from each of the two models. It is evident that they all are highly correlated. Nonetheless, the formal test for equality between the option-implied and the nonparametric diffusive volatility estimates rejects the null hypothesis for a nontrivial number of days for all models, as may be confirmed from Panel C of Table XI. In fact, the rates are fairly similar across the two models, likely reflecting the tighter standard errors on the option-implied volatility estimates associated with the three-factor model. This is corroborated by the serial correlation in the discrepancy between the option implied and nonparametric volatility estimates plotted in the bottom panels of Figure 7. Under correct model specification, these series should not display significant autocorrelation. However, relatively strong temporal dependence is evident for both series, albeit somewhat less for the three-factor model.

TABLE XII
PARAMETER STABILITY TESTS ON S&P 500 OPTIONS DATA^a

Parameter	Nominal Size of Test		Parameter	Nominal Size of Test	
	1%	5%		1%	5%
Panel A: One-Factor Model					
ρ_d	25.71%	40.00%	λ_j	68.57%	75.24%
\bar{v}	72.38%	77.14%	μ_x	47.62%	58.10%
κ	81.90%	86.67%	σ_x	30.48%	38.10%
σ_d	69.52%	74.29%	μ_v	41.90%	50.48%
			ρ_j	32.38%	36.19%
Panel B: Three-Factor Model					
ρ_1	0.00	0.00	ρ_3	0.00	0.00
\bar{v}_1	49.52	57.14	c_0^+	18.09	26.67
κ_1	47.62	56.19	c_1^-	12.38	18.10
σ_1	25.71	34.29	c_1^+	0.00	2.85
ρ_2	0.00	0.00	c_2^-	0.00	0.00
\bar{v}_2	34.29	43.81	c_2^+	0.00	0.00
κ_2	73.33	82.86	c_3^-	13.33	17.14
σ_2	13.33	25.91	λ_-	35.24	45.71
μ_u	0.00	0.00	λ_+	6.67	8.57
κ_3	32.38	42.86	μ_1	7.62	15.24

^aTests based on parameter estimates of the models over consecutive calendar years in the sample. The test is based on Corollary 2.

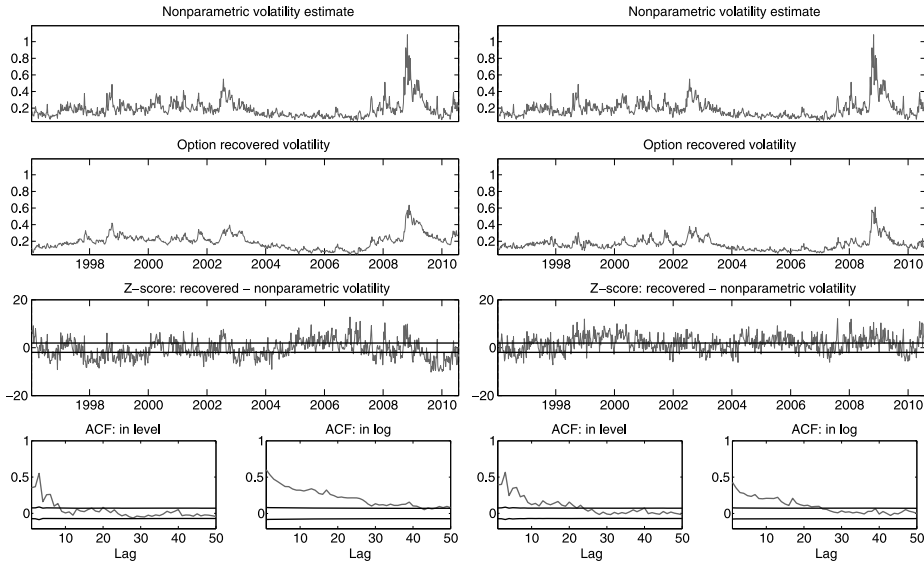


FIGURE 7.—Volatility estimates. The left panel corresponds to the one-factor model and the right panel to the three-factor model. The bottom plots of the figure are the autocorrelations in $\xi_1(\hat{S}_t) - \hat{V}_t^n$ and $\log(\xi_1(\hat{S}_t)) - \log(\hat{V}_t^n)$.

S2.2. Implied Volatility Skews for the Three-Factor Model

We now provide implied volatility skews for each calendar year for the three-factor model (22) (see Figures 8–22).




FIGURE 8.—Implied volatility standard error bands, January 1, 1996–January 1, 1997. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 9.—Implied volatility standard error bands, January 1, 1997–January 1, 1998. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 10.—Implied volatility standard error bands, January 1, 1998–January 1, 1999. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 11.—Implied volatility standard error bands, January 1, 1999–January 1, 2000. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 12.—Implied volatility standard error bands, January 1, 2000–January 1, 2001. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 13.—Implied volatility standard error bands, January 1, 2001–January 1, 2002. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 14.—Implied volatility standard error bands, January 1, 2002–January 1, 2003. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 15.—Implied volatility standard error bands, January 1, 2003–January 1, 2004. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 16.—Implied volatility standard error bands, January 1, 2004–January 1, 2005. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 17.—Implied volatility standard error bands, January 1, 2005–January 1, 2006. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 18.—Implied volatility standard error bands, January 1, 2006–January 1, 2007. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 19.—Implied volatility standard error bands, January 1, 2007–January 1, 2008. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 20.—Implied volatility standard error bands, January 1, 2008–January 1, 2009. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 21.—Implied volatility standard error bands, January 1, 2009–January 1, 2010. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).




FIGURE 22.—Implied volatility standard error bands, January 1, 2010–July 21, 2010. Left panel: short-maturities (tenor below 60 days). Right panel: long-maturities (tenor exceeds 150 days).

S2.3. Alternative Three-Factor Volatility Model

We conclude this section with results from estimation of an alternative three-factor model. It stipulates the following risk-neutral equity index dynamics:

$$\begin{aligned}
 (73) \quad \frac{dX_t}{X_{t-}} &= (r_t - \delta_t) dt + \sqrt{V_{1,t}} dW_{1,t} + \sqrt{V_{2,t}} dW_{2,t} + \sqrt{V_{3,t}} dW_{3,t} \\
 &\quad + \int_{\mathbb{R}^2} (e^x - 1) \tilde{\mu}(dt, dx, dy), \\
 dV_{1,t} &= \kappa_1(\bar{v}_1 - V_{1,t}) dt + \sigma_1 \sqrt{V_{1,t}} dB_{1,t} \\
 &\quad + \mu_{v_1} \int_{\mathbb{R}^2} x^2 1_{\{x < 0\}} \mu(dt, dx, dy), \\
 dV_{2,t} &= \kappa_2(\bar{v}_2 - V_{2,t}) dt + \sigma_2 \sqrt{V_{2,t}} dB_{2,t}, \\
 dV_{3,t} &= -\kappa_3 V_{3,t} dt + \mu_{v_3} \int_{\mathbb{R}^2} [(1 - \rho_3)x^2 1_{\{x < 0\}} + \rho_3 y^2] \mu(dt, dx, dy).
 \end{aligned}$$

The jump measure μ has a compensator given by $dt \otimes \nu_t^{\mathbb{Q}}(dx, dy)$, where

$$\begin{aligned}
 \nu_t^{\mathbb{Q}}(dx, dy) &= \left\{ (c^- 1_{\{x < 0\}} \lambda_- e^{-\lambda_- |x|} + c^+ 1_{\{x > 0\}} \lambda_+ e^{-\lambda_+ x}) 1_{\{y=0\}} \right. \\
 &\quad \left. + c^- 1_{\{x=0, y < 0\}} \lambda_- e^{-\lambda_- |y|} \right\} dx \otimes dy,
 \end{aligned}$$

$$c^- = c_0^- + c_1^- V_{1,t-} + c_2^- V_{2,t-} + c_3^- V_{3,t-},$$

$$c^+ = c_0^+ + c_1^+ V_{1,t-} + c_2^+ V_{2,t-} + c_3^+ V_{3,t-}.$$

TABLE XIII
PARAMETER ESTIMATES FOR THE ALTERNATIVE THREE-FACTOR MODEL^a

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	-0.8225	0.0195	σ_2	0.0838	0.0076	c_2^-	79.3262	3.7170
\bar{v}_1	0.0127	0.0003	μ_{v_3}	2.9428	0.962	c_2^+	247.1285	99.4715
κ_1	5.7246	0.1274	κ_3	19.9491	0.9336	c_3^-	123.8232	9.7126
σ_1	0.3587	0.0101	ρ_3	0.8109	0.2020	λ_-	14.7842	0.5526
ρ_2	-0.9964	0.0569	c_0^+	1.9995	1.3937	λ_+	102.3357	3.4339
\bar{v}_2	0.0318	0.0043	c_1^-	0.1289	1.1831	μ_{v_1}	10.4981	1.0545
κ_2	0.1688	0.0249	c_1^+	290.4399	78.9567			

^aParameter estimates of the alternative three-factor model (73) for S&P 500 equity-index option data sampled every Wednesday over the period January 1996–July 2010.

The specification in (73) differs from our original three-factor model (22) primarily by having the third factor, driving the jump intensity, V_3 , be a component of the diffusive volatility. The number of parameters in the two alternative three-factor models is identical.

As for model (22), we impose $c_3^+ = 0$ and $c_0^- = 0$ during the estimation. The parameter estimates of model (73) are reported in Table XIII, the corresponding Z-scores for the fit to the separate regions of the option surface are plotted on Figure 23, and results for further diagnostic tests and plots are given in Table XIV and Figure 24.

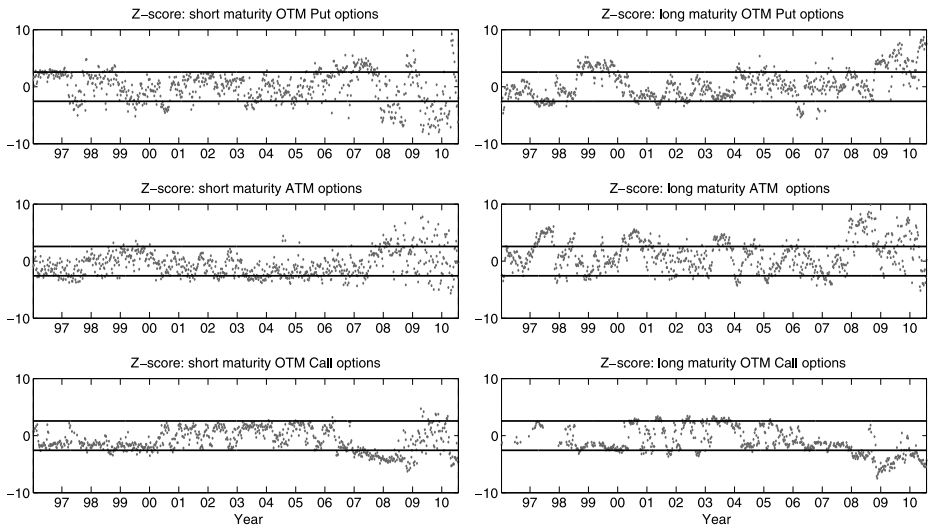


FIGURE 23.—Option price fit for the alternative three-factor model.

TABLE XIV
DIAGNOSTIC TESTS FOR THE ALTERNATIVE THREE-FACTOR MODEL^a

Test	Alternative Three-Factor Model Nominal Size	
	1%	5%
Panel A: Fit to the Option Panel		
OTM, short-maturity puts	34.08%	49.61%
ATM, short-maturity puts	24.61%	42.37%
OTM, short-maturity calls	20.53%	38.55%
OTM, long-maturity puts	30.26%	41.84%
ATM, long-maturity puts	38.82%	51.05%
OTM, long-maturity calls	26.84%	47.37%
Panel B: Root-Mean-Squared-Error of IV Option Fit 1.66%		
Panel C: Equality of Implied and Nonparametric Volatility 43.94% 55.52%		

^aPanel A reports rejection frequencies across the full sample for the option fit to specific portions of the option surface at the end of trading on Wednesdays. This test is based on Corollary 1, using the first two maturities for the three initial tests and all remaining options with maturity less than one year for the last three tests. DOTM puts, OTM puts and calls, and short- versus long-maturity options are defined in Figure 1. Panel B provides the root-mean-squared-error of the model-implied BSIV relative to the market mid-quote BSIV across all options used for estimation over the full sample. The test in Panel C is defined in Corollary 3.

Finally, in Figure 25, we provide a more direct comparison of the two three-factor models in terms of the distribution of the Z -scores relative to the theoretical quantiles. We find that the extended model with an independent degree of flexibility in the left jump tail produces quantiles that are positioned much closer to the theoretical 45 degree benchmark for five of the six regions relative to the traditional three-factor model, while they are very similar for the short-maturity calls. Overall, the model constructed according to the familiar volatility structure falls significantly short of the new extended three-factor model along the majority of the dimensions explored, and often by a substantial margin.

S3. DETAILS ON COMPUTATION

The estimation procedure described in Section 4.2 entails two distinct issues. On the one hand, on each day in the sample, we have to recover the volatility states inverting the option pricing formula for a given model for the asset returns dynamics. We use the free/open-source NLOpt⁴¹ library for nonlinear optimization to perform this task. Specifically, we employ the BOBYQA

⁴¹For further information about the library and the different minimization algorithms, see the official web-site http://ab-initio.mit.edu/wiki/index.php/NLOpt_Introduction.

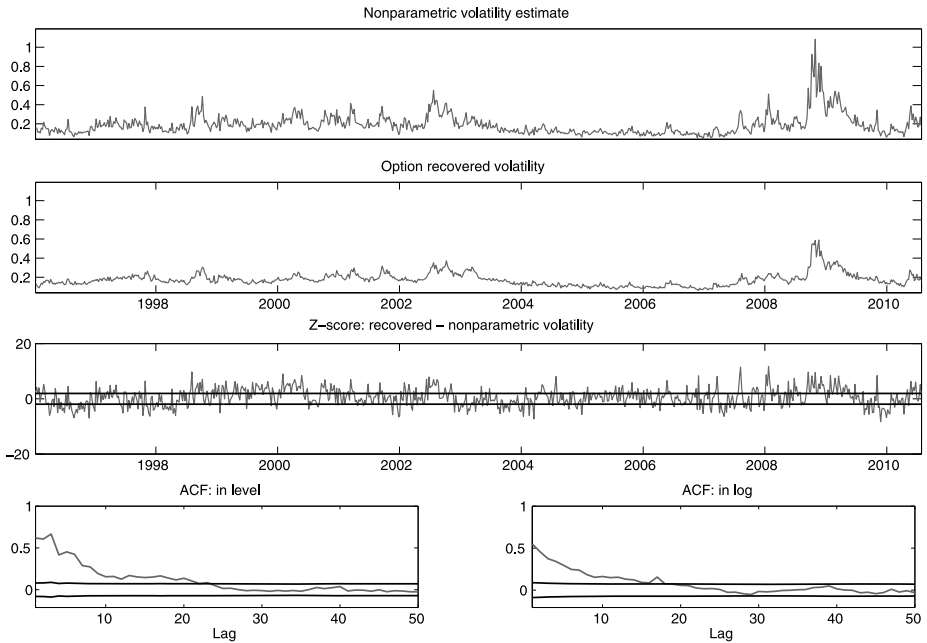


FIGURE 24.—Volatility estimates for the alternative three-factor model. The bottom plots of the figure are the autocorrelations in $\xi_1(\widehat{S}_t) - \widehat{V}_t^n$ and $\log(\xi_1(\widehat{S}_t)) - \log(\widehat{V}_t^n)$.

algorithm: this is a local derivative-free optimization algorithm which performs derivative-free bound constrained optimization using an iteratively constructed quadratic approximation for the objective function. This method proved to be fast and reliable.

On the other hand, we need to minimize the objective function in equation (5) which depends upon a high-dimensional parameter vector and it is costly to evaluate since it nests many minimizations coming from the volatility state recovery that has to be done for every day in the sample (i.e., about three thousand observations). To overcome this problem, we followed four complementary strategies. First, all the code has been written in C++ to benefit from its computational speed. Second, since the inversion problem is inherently independent from one day to another, we relied on the Open MPI (Messages Passing Interface, www.open-mpi.org) library in order to exploit the power of multiple CPUs at the same time, which means that we can simultaneously back out the volatility states over different days. Third, we chose the Fourier-cosine series expansion described in Fang and Oosterlee (2009) as our option valuation method, which has been shown to be remarkably faster than the Carr and Madan (1999) method. Fourth, since the Fourier-cosine series expansion method basically relies on the knowledge of the log-price characteristic function (CF), we needed a way to compute it quickly even when it is not known

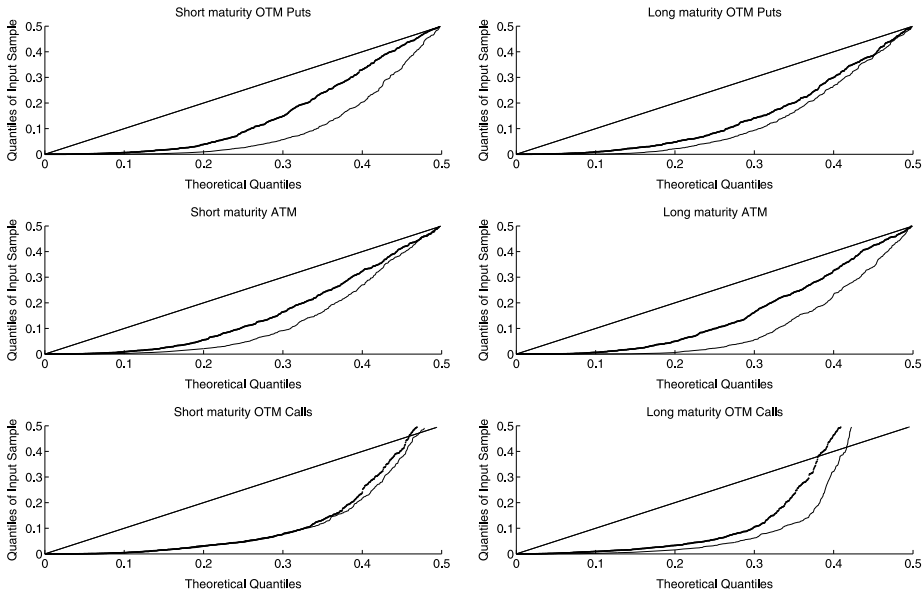


FIGURE 25.—The dark line corresponds to our three-factor model and the gray line to the three-factor volatility model.

in closed form. This is easily done taking into account that the CF for example for our two-factor volatility model is of the form:

$$\begin{aligned}
 f(\tau, y_t, v_{1,t}, v_{2,t}, u) &= \mathbb{E}_t[e^{x_T u}] \\
 &:= e^{\alpha(\tau, \theta, u) + \beta_1(\tau, \theta, u)v_{1,t} + \beta_2(\tau, \theta, u)v_{2,t} + ux_t}, \quad u \in \mathbb{C},
 \end{aligned}$$

where $x_t = \log(X_t)$. The coefficients $\alpha(\tau, \theta, u)$ and $\beta_i(\tau, \theta, u)$ can be computed once at the beginning of each objective function evaluation. In this way, we only need to solve the system of ODEs over the longest option maturity in our sample and for different values of u once, for each parameter vector.

Finally, in order to cope with time constraints, we followed two different approaches for the Monte Carlo study and the empirical investigation, respectively. Precisely:

- Monte Carlo study: we carried out the minimization using the NLOpt library. Specifically, we sequentially used a global search algorithm and a local search one. We started the minimization from the true parameter value but we allowed a wide exploration of the parameter space through the global search algorithm. We used the Controlled Random Search (CRS) with local mutation algorithm as our global optimization: it can be compared to genetic algorithms, since it starts with a random “population” of points and then it randomly evolves them. Finally, the local search has been done with the Sbpplx

(based on Subplex) algorithm, which has been proven to be more efficient and robust than standard simplex methods.

- Empirical application: the Markov chain Monte Carlo (MCMC) method has been used to perform the objective function minimization. We employed the wide-scope C++ library of Ronald Gallant which is an implementation of the Chernozhukov and Hong (2003) estimator and which can be downloaded from <http://www.unc.edu/~arg/>.

In order to carry out the Monte Carlo study in a timely fashion, we relied on a High Performance Computing System—ranked among the 500 fastest computers worldwide—composed of 504 (4032 cores) Intel Nehalem E5520, 64-bit, 8 M Cache, 2.26 GHz, with 48 GB's of DDR3 memory per node, and 252 (3024 cores) Intel Westmere X5650, 64-bit, 12 MB Cache, 2.66 GHz, with 48 GB's of QDR memory per node. The complete experiment required approximatively 100 thousand CPU hours.

REFERENCES

- CARR, P., AND D. MADAN (1999): “Option Valuation Using the Fast Fourier Transform,” *Journal of Computational Finance*, 2, 61–73. [16]
- CHERNOZHUKOV, V., AND H. HONG (2003): “An MCMC Approach to Classical Estimation,” *Journal of Econometrics*, 115, 293–346. [18]
- FANG, F., AND C. OOSTERLEE (2009): “A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions,” *SIAM Journal on Scientific Computing*, 31, 826–848. [16]

Dept. of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, U.S.A., NBER, Cambridge, MA, U.S.A., and CREATES, Aarhus, Denmark; t-andersen@northwestern.edu,

Carey Business School, The John Hopkins University, Baltimore, MD 21202, U.S.A.; nicola.fusari@jhu.edu,

and

Dept. of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, U.S.A.; v-todorov@northwestern.edu.

Manuscript received April, 2012; final revision received December, 2014.